

**AFRL-IF-WP-TR-2005-1505**

**POWER MANAGEMENT FOR REAL-  
TIME SYSTEMS**

**Rami Melhem, Daniel Mosse, Mootaz Elnozahy, and  
Hazim Shafi**

**The University of Pittsburgh  
Department of Computer Science  
Pittsburgh, PA 15260**

**JULY 2004**

**Final Report for 04 August 2000 – 30 June 2004**



**Approved for public release; distribution unlimited.**

**STINFO FINAL REPORT**

**INFORMATION DIRECTORATE  
AIR FORCE RESEARCH LABORATORY  
AIR FORCE MATERIEL COMMAND  
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7334**

# NOTICE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report has been reviewed and is releasable to the National Technical Information Service (NTIS). It will be available to the general public, including foreign nationals.

THIS TECHNICAL REPORT HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION.

\_\_\_\_\_/s/\_\_\_\_\_  
JONG S. HWANG, Program Monitor

\_\_\_\_\_/s/\_\_\_\_\_  
ROBERT A. EHRET, Chief  
Collaborative Simulation Technology & Applications Branch  
Information Systems Division  
Information Directorate

This report is published in the interest of scientific and technical information exchange and does not constitute approval or disapproval of its ideas or findings.

<b>REPORT DOCUMENTATION PAGE</b>					<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. <b>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</b>						
<b>1. REPORT DATE (DD-MM-YY)</b> July 2004		<b>2. REPORT TYPE</b> Final		<b>3. DATES COVERED (From - To)</b> 8/4/00 – 6/30/04		
<b>4. TITLE AND SUBTITLE</b> POWER MANAGEMENT FOR REAL-TIME SYSTEMS				<b>5a. CONTRACT NUMBER</b> F33615-00-C-1736		
				<b>5b. GRANT NUMBER</b>		
				<b>5c. PROGRAM ELEMENT NUMBER</b> 69199F		
<b>6. AUTHOR(S)</b> Rami Melhem, Daniel Mosse, Mootaz Elnozahy, and Hazim Shafi				<b>5d. PROJECT NUMBER</b> ARPI		
				<b>5e. TASK NUMBER</b> FS		
				<b>5f. WORK UNIT NUMBER</b> AA		
<b>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)</b>  The University of Pittsburgh Department of Computer Science Pittsburgh, PA 15260				<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b>		
<b>9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b>  Information Directorate Air Force Research Laboratory Air Force Materiel Command Wright-Patterson AFB, OH 45433-7334  Defense Advanced Research Projects Agency 3701 N. Fairfax Drive Arlington, Virginia 22203-1714				<b>10. SPONSORING/MONITORING AGENCY ACRONYM(S)</b> AFRL/IFSD		
				<b>11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S)</b> AFRL-IF-WP-TR-2005-1505		
<b>12. DISTRIBUTION/AVAILABILITY STATEMENT</b> Approved for public release; distribution is unlimited.						
<b>13. SUPPLEMENTARY NOTES</b>						
<b>14. ABSTRACT</b> The Power Aware Real-Time Scheduling (PARTS) project is based on attacking the power management problem for real-time systems at the system level. This includes both modifications to the applications, to the operating system, as well as compiler modifications to insert so-called power management points. The PARTS project started with scheduling independent tasks in a single CPU and has expanded to multiple CPUs as well as tasks with dependencies. We approached the power management problem from three different angles. The first is the minimization of energy consumption for given timing constraints. The second is the maximization of the system's reward (utility) for specific timing and energy restrictions (including rechargeable systems). The third is the tradeoff between energy consumption and reliability for specific timing and performance restrictions. Further, we validated the theory developed by carrying out implementations on the Power Aware Multiprocessor Architecture platform (PAMA), which consists of four rad-hard PowerPC 750 processors running actual space applications.						
<b>15. SUBJECT TERMS</b> Real-Time Systems, Deadline, Power Management, Critical Systems						
<b>16. SECURITY CLASSIFICATION OF:</b>			<b>17. LIMITATION OF ABSTRACT:</b> SAR	<b>18. NUMBER OF PAGES</b> 20	<b>19a. NAME OF RESPONSIBLE PERSON (Monitor)</b> Jong S. Hwang <b>19b. TELEPHONE NUMBER (Include Area Code)</b> (937) 904-9048	
<b>a. REPORT</b> Unclassified	<b>b. ABSTRACT</b> Unclassified	<b>c. THIS PAGE</b> Unclassified				

## Table of Contents

Abstract .....	iv
1. Introduction .....	1
2. Power management points .....	1
3. Dynamic voltage scaling in single and multiple processors .....	2
4. Reward-based scheduling .....	2
5. Fault tolerance .....	3
6. Power management beyond the CPU .....	3
7. Implementation and simulation .....	4
8. Technology Transition and community activities .....	6
9. Publications resulted from the project .....	8

## **Abstract**

The Power Aware Real-Time Scheduling (PARTS) project is based on attacking the power management problem for real-time systems at the system level. This includes both modifications to the applications, to the operating system as well as compiler modifications to insert so-called power management points. Most of the power management research is based on reducing the CPU voltage and thus obtaining a quadratic gain in energy savings with linear increase in time. The PARTS project started with scheduling independent tasks in a single CPU and has expanded to multiple CPUs as well as tasks with dependencies. We approached the power management problem from three different angles. The first is the minimization of energy consumption for given timing constraints. The second is the maximization of the system's reward (utility) for specific timing and energy restrictions (including rechargeable systems). The third is the tradeoff between energy consumption and reliability for specific timing and performance restrictions. Further, we validated the theory developed by carrying out implementations on the Power Aware Multiprocessor Architecture platform (PAMA), which consists of four rad-hard PowerPC 750 processors running actual space applications.

## 1. Introduction

The goal of this project was to develop and implement techniques for power management in real-time systems, where applications have to adhere to strict deadlines. During the course of this project, the PIs and their students studied the tradeoffs involved when the slack that exists, or becomes available dynamically, in the system is used for one of three goals: (a) reducing the energy consumption, (b) improving the system's performance, and (c) increasing the reliability. The schemes we devised rely on different management techniques including scheduling algorithms, speed control of CPUs, and dynamic power monitoring and mode changes. We integrated our new schemes into the appropriate components of the system (e.g., the kernel scheduler or a library). We developed such schemes for single and multiprocessor systems, for tasks with and without precedence constraints and for deadline-driven as well as rate-based systems. We applied these schemes to the PAMA subscale demonstration which is a multiprocessor platform being built by BAE Systems and ISI East, and we evaluated the resulting energy savings. Our scheduling schemes were also tested on real-life software, which came from our industrial partners in Phase I and II of the Power Aware Computing and Communication program (PAC/C). These applications are Automated Target Recognition (ATR), Event Extraction (EE), and Complex Ambiguity Function (CAF).

As part of our project, we also created a cycle-accurate and power-accurate emulator for the PowerPC 405LP and the PowerPC 750 to measure power consumption of each of the resources; the focus was mainly in the CPU at the lowest-possible level (i.e., instruction level). This new power-aware emulator provides an interface to the higher layers of our proposed architecture for querying and changing specific power parameters of the system; it is also possible to boot the Linux operating system on top of the emulator, allowing for multiple processes to be run and studied.

In the remainder of this report, we describe, in some details, the accomplishments achieved during the course of the project.

## 2. Power Management Points (PMP)

We developed a theoretical formulation to estimate the granularity of the frequency of invoking power management points (PMPs)[31]. A PMP is essentially a piece of code that carries out speed (voltage and frequency) changes. This estimation takes into consideration the tradeoff between the need of frequent speed changes for better slack (idle time) management and the overhead of speed changes. We introduced a collaborative approach between the compiler and the operating system (OS) that uses fine-grained information about the execution times of a real-time application to reduce energy consumption [21]. Specifically, the compiler annotates an application's source code with path-dependent information called power management hints (PMHs)[20]. This information captures the temporal behavior of the application, which varies when executing different paths. During program execution, the OS periodically changes the processor's frequency and voltage based on the temporal information provided by the PMHs. We evaluated our scheme using four embedded applications: a video decoder, automatic target recognition, audio encoder, and subband filter event extraction. Our scheme shows an energy reduction of up to 90% over *no power management* and up to 50% over *static power management* schemes. As an example, our scheme, when tested on the ATR code obtained from Northrop Grumman, shows that energy reduction ranges from 10% to 90% over no-power-management for the ATR benchmark provided, depending on the required frame processing frequency (mode of operation) and the number of detected potential targets.

### 3. Dynamic Voltage Scaling (DVS) in single and multiple processors

One of the open areas that we provided solutions for is DVS for periodic hard real-time tasks. Our solution includes three parts: (a) a static, off-line solution to compute the optimal speed, assuming worst-case workload for each arrival, (b) an on-line speed adjustment mechanism to reclaim energy not used by tasks that complete early, and (c) an on-line, adaptive and speculative speed adjustment mechanism to anticipate early completions of future executions. The speculative scheme makes use of the average-case workload information to save energy, while still guaranteeing all deadlines. Our results show that the reclaiming algorithm saves a striking 50% of the energy over the static algorithm. Further, we experimented with more aggressive speculative techniques and showed that up to an additional 20% savings can be obtained over non-aggressive reclaiming algorithms.

We have also studied power management for rate-based soft real-time systems, in which there is a rate requirement but no hard deadlines to be obeyed (e.g., Web servers, EE and CAF). By analyzing the effect of the static power on the power management scheme, we have devised several methods for determining when voltage scaling is applicable and when turning servers ON/OFF is more appropriate, with respect to energy savings, or the combination thereof. We have also devised a scheme for determining the optimal system speed to minimize the energy consumption for any application while respecting the application's timing constraints; this scheme is based on the static/dynamic power ratio of the system, which can be determined *a priori*[15].

Power management in single-processors was extended to parallel and distributed system. As a basis for the power management in parallel systems, we developed a novel scheduling algorithm called *boundary fair scheduling* for periodic hard real-time tasks executing on multiprocessor systems[19]. This algorithm is optimal in the sense of achieving full system utilization and can be modified to manage power for periodic tasks on parallel systems. We have also developed static power management schemes for parallel and distributed frame-based systems. These schemes consider the degree of parallelism of the applications when allocating global static slack. In addition to static power management, a dynamic power management scheme was introduced for distributed systems, in which we not only take into account the communication between tasks, but also optimize the schedules based on the gaps that are created due to communication and precedence constraints[22]. A detailed study of power management for shared memory multiprocessor systems were conducted assuming a general and/or processor graph model and different slack sharing techniques were designed to guarantee correctness and efficiency of execution[7,11,27,29]. Power savings of up to 65% over non-slack sharing schemes were obtained in simulation studies for the ATR, EE and CAF benchmarks.

### 4. Reward-Based Scheduling

In addition to systems that have to execute specific sets of tasks, we considered more general systems in which a tradeoff exists between the cost (with respect to resources) and the benefit (with respect to utility) of executing each application for a specific amount of time[1,12,16]. In such systems, different applications have different rewards/values and within a single application there might be different levels of accuracy (depending on how much time the application uses the CPU). Our work combines three dimensions, namely—timeliness, energy, and rewards—with the goal of maximizing the system reward for energy-constrained real-time systems. We developed algorithms to solve a number of optimization problems involving the above tradeoffs[6,13]. To our knowledge, ours is the first solution that combines the three constraints mentioned above. We formalized the problem of maximizing the rewards while meeting the deadline without exceeding a given energy budget. We showed that if the power functions of all the tasks have the same form

(for example, they are all quadratic or cubic on the processor speed), then reward is maximized when the processor power is the same for all tasks and we devised a closed-form solution to the problem, which had been an open problem until then.

For the more general case of power functions with different forms, we have proposed an iterative algorithm that proved to be extremely fast and accurate[24]. The algorithm is based on repeatedly trading reward for energy, within the timing constraints, until the energy constraints are satisfied. These algorithms run in only a few microseconds even for tens of tasks, and thus are suitable for real-time and embedded systems.

After laying the ground work with the tri-valued approach above, we have extended our schemes to systems in which each application may have multiple versions, adding yet another degree of freedom[23]. We have further included in the model the requirement that devices may harvest energy from the environment (typically there is a cycle of discharging and recharging). Our results again show that most of the algorithms are extremely fast (run in few microseconds) and accurate (within 1-2% of the optimal performance).

## 5. Fault Tolerance

We obtained results related to the issue of reliability in power aware real-time systems. We have compared Duplex and Triple Modular Redundancy (TMR) systems, both from an analytical modeling perspective, as well as quantitative perspective, and we were able to characterize the system parameters in which each technique is superior to the others[10,25]. Our characterization allows the user and designer to focus on other aspects in the system design, such as fault coverage, degree of reliability required, volume, or cost.

In our analysis we designed a new type of systems, namely *optimistic TMR* systems, in which we turn off one component of the TMR system whenever possible. Should a failure be detected by the two remaining components, the third system is brought out from hibernation and is allowed to execute as fast as possible to produce the tie-breaking vote before the deadline[15,25]. Optimistic TMR systems are more efficient, in terms of energy consumption, than power-aware Duplex and TMR systems. The analysis of optimistic-TMR shows up to 20% saving in energy over simple TMR without any deterioration in reliability.

Our analysis technique allows systems integrators to determine the most energy-efficient fault-tolerant structures to be used in a system/application for a given system configuration (including load, applications, hardware, etc). The system integrator can choose from *duplex*, *TMR*, *Optimistic TMR*, *checkpointing*, and a *hybrid checkpointing-redundancy* schemes. We extended the analysis for parallel real-time applications running on a system consisting of a fixed number of processing units, taking into account several parallel recovery schemes. For example, we have combined modular redundancy with check-pointing, and developed a methodology to find the most energy efficient configuration to tolerate a certain number of transient faults.

## 6. Power Management beyond the CPU

We extended our study to investigate the effect of combining CPU power management schemes with other energy management techniques, such as those used to manage the power in the wireless communication system, memory and hard disks, with the goal of improving the overall consumption of the entire system. Energy savings can be achieved in memory by setting unused memory banks to a lower power level to consume less energy and in hard disks by stop spinning during idle times.



Beyond wired distributed systems, we started analyzing the next natural extension of the power management schemes, namely wireless networks. One of the constraints for building an efficient adhoc network is finite battery supply. Some previous work proposed the idea of minimizing the transmission power by sending the data in a multi-hop fashion to the destination by relaying the packets at intermediate closer nodes. We devised a *unified* model that considers the energy wasted in collisions and the interference level of neighboring nodes together (the interference/collision combination has been largely ignored in the past). From this model we derived the total network throughput and the total energy consumption in the network and proved that it is not always optimal to use the minimum transmission power. We reported a throughput increase of 30% with the correct choice of transmission power over the minimum transmission power without a significant increase in the total network energy consumption[18]. We have also developed a model of rate/energy comparison, to determine the energy-efficient bandwidth at which wireless nodes should send data packets. This scheme takes into account the battery levels of nodes in a path and determines the rate at which nodes in a path should send and receive data packets. We are currently using dynamic source routing (DSR) as our routing algorithm, but the scheme is general enough to be readily extended to other routing algorithms[17].

We also worked on the problem of designing a new energy-efficient MAC protocol for adhoc networks; we were able to reduce the network energy consumption and to extend the network lifetime by decreasing the number of collisions faced in the network. This is done through a dynamic priority scheme that is based on the node's current residual energy. Our new MAC protocol achieves a 40% increase in the network lifetime and is also backward compatible with the widely deployed 802.11 DCF protocol.

For memory, we wired the Pecan prototype (an evaluation system based on the 405GP with 9 power planes, for power measurement of portions of the system) to allow the placement of memory banks in self-refresh mode through software control. When placed in this mode, the memory subsystem consumes approximately 3% of active power. We also modified the Linux operating system to subdivide memory into two regions. One region operates normally, while the second one operates as a RAM disk. The RAM disk is put in self-refresh mode most of the time, except when the operating system reads or writes into the RAM disk under software control. The resulting savings were on the order of 3W (about 50% of the total power consumed by the board). This experiment shows that intelligent memory management can result in significant power savings for ultra-low power systems.

We developed a new simulator, MemSim, to model joint power-performance characteristics of main memory systems; the simulator allows us to model the multiple datapaths, frequency domains, complex memory organizations, modern server memory (SDRAM DDR) page and power modes accurately. In addition, we developed the MEMPOWER toolset for memory power analysis. MEMPOWER takes as input memory reference traces and calculates the memory's power, energy and additional delay, if any, under various power management policies. It also provides a detailed analysis of the inter-arrival patterns and spatial distribution of references in the trace. Since most policies for memory power management are threshold based, MEMPOWER includes a set of programs that compute the optimal threshold value for the reference pattern found in the trace.

## 7. Implementation and Simulation

A major part of our simulations were performed on SimpleScalar, which is a low-level simulator for RISC architectures, and on Wattch, which extends SimpleScalar to estimate the CPU power consumption. The variable-speed SimpleScalar allows us to test the performance of a variable-voltage architecture and to adjust the CPU speed to meet a certain quality of service. This

variable-voltage enhanced SimpleScalar was initially the primary tool used by our group in evaluating architectural changes needed in the processors for improving the energy consumption of systems.

We conducted a series of power measurements for various workloads and synthetic benchmarks, using a range of processors, including the Transmeta Crusoe TM5400 chip, the PowerPC 750, and a conventional Intel 700MHz system. Results show that it will be important to implement our work on processors that have aggressive power management features such as the Crusoe[37].

During the project, instead of relying on the inaccuracies of Wattch, we developed an event and execution-driven simulator, called CycleSim, which is part of the MAMBO simulation environment at IBM. CycleSim models an IBM PowerPC 405GP system-on-a-chip and includes a cycle-accurate, event-based power model of the processor core and enough system details to boot an operating system[39]. We created the event-based power model for the 405 core by measuring the cost of various architectural events such as cache hits or misses, table lookaside buffer (TLB) hits or misses, various instruction types, etc. These measurements were performed on the Pecan board (an evaluation system based on the 405GP with 9 power planes, for power measurement of portions of the system) that was designed and implemented with DARPA support. When the simulator was completed, we performed timing and power validation of the simulator against the actual hardware using 42 EEMBC (Embedded Microprocessor Benchmarks Consortium) applications. Preliminary results using four applications (text, iDCT, matrix, and Viterbi) show that our timing simulation is within 1% (with a standard deviation of 2.5%) and the average error in power prediction was less than 5% (with a standard deviation of 5.1%) compared with the actual hardware.

We extended MAMBO (and then the design of the 405GP) to implement the voltage and frequency scaling capabilities of the PPC 405LP. Specifically, IBM provided a source-level license agreement to the University of Pittsburgh, and supported a Pitt student as a summer intern to make power measurements on the PowerPC 405LP processor. We then used those measurements to upgrade the CycleSim simulator to support dynamic voltage and frequency scaling. Moreover, we augmented MAMBO with support for file system calls to read and write multiple files, and with the capability of handling on-line, real-time input, memory banks with power consumption model, and standby/sleep modes for memory banks. Finally, we also created software tools to aid in the analysis of the results provided by the MAMBO emulator.

### **Porting Linux to MAMBO**

Our work with MAMBO had initially been done without any operating systems, and in the past we have ported the SimOS emulator to an x86/Linux based platform. We have integrated Linux with the MAMBO simulator, to realize a flexible simulation environment that offers a tradeoff between accuracy and speed of execution. This work allowed fine-grained processor measurements and emulation of power consumption to be integrated with a system-wide power measurement tool.

### **Enhancing Linux with DVS**

We also worked at the kernel level, implementing all the necessary system calls and supporting libraries to enable applications and schedulers to be DVS-aware. This entailed simple kernel modifications that can be easily ported to Linux systems through kernel patches. We demonstrated such capability in 2002 during a PI meeting, with a GUI that allowed a console root user to change the speed of specific tasks.

## Singularity Board

We developed a prototype power-aware dense storage cache blade, dubbed *Singularity*. We have completed board verification, power-on-self-test and bring-up of both the back-end NPe405L sub-systems and the front-end *network processor*. The Linux kernel now boots on the back-end processors and the IBM Advanced Software Offering (Layer 4 Switching) boots on the Network Processor.

In addition to the verification and bring-up process, we have been doing background research into payload caching, alternative storage network protocols, and collaborative page caching. This enabling research will allow IBM to get baseline benchmarks on the platform and to begin implementing various power management, request distribution, and collaborative caching schemes using the Singularity board.

## 8. Technology Transition and Community Activities

We worked extensively with BAE Systems to apply our voltage scaling schemes to space applications, in general, and to the Subband Tuner (EE) and Complex Ambiguity Function (CAF) applications, in particular. Upon a request from BAE Systems, we designed a simple power management interface that can be used by application developers to allow them to benefit from power-management technology. This interface was well received, enhancing the PAMA application interface, in the form of a library of calls that can be used by any application.

We will continue to work with BAE Systems in Phase II of the PAC/C program. In addition to transferring our power management technology to the AMPS prototype, we are developing a cycle-accurate model of the PowerPC-based RAD750 microprocessor. We also created a power model for the processor using our event-based methodology and delivered the simulator to BAE Systems in support of their efforts. We are collaborating with other research groups to enhance our power simulation infrastructure and use it to perform research on power-aware operating systems, compilers, and memory subsystem design.

We added support to Linux for throttling features present in the PowerPC 970 and characterized the performance and power behavior of these features. We enhanced Linux power measurement for the PowerMac G5 and we characterized the behavior of the PowerMac's built-in power measurement capabilities.

Following up on our work to prototype support for turning off portions of the main memory on the Pecan board, we contacted several organizations in IBM on how to build power-aware memory controllers that can use the firmware and software that we distributed. IBM has decided to showcase this work in the Embedded Systems Conference in March 2004 to boost the marketability of the work.

We busily promoted the use of the MAMBO simulator and our power modeling techniques in both industry and academia. Towards this goal, IBM has signed a MAMBO license agreement with Vanderbilt University (also funded by PAC/C) and assisted their efforts in supporting MAMBO in their MILAN simulation environment. In addition, we presented a MAMBO tutorial during ISPASS (The IEEE International Symposium on Performance Analysis of Systems and Software) in March 2003 that included details on the design and validation of our power modeling efforts. The tutorial was well attended and received. More recently, the MemSim simulator was licensed to Penn. State University to help with memory power related research.

The PARTS team members are well-known players in the power management community, participating in many of the organization committees of well-recognized conferences throughout the USA and the world. This enhances the visibility of the project's accomplishments and allows

these to be transitioned to industry. In the real-time research community, power-management sessions have become a feature in all conferences, whereas three years ago (before the start of the PAC/C program) no paper was presented on this issue.

The PIs have been involved in the dissemination of power-aware technologies, through organizing special sessions and workshops on power aware computing in real-time conferences and the publication of a book "Power Aware Computing", edited by Rami Melhem and Robert Graybill, and published by Kluwer/Plenum, mostly with PAC/C researchers as authors. In the real-time research community, it is surprising to realize that a full 10-12% of the papers submitted to RTSS and RTAS, the premier real-time conferences, deal with power-aware systems.

In the same vein, we organized the ACEED (Austin Conference for Energy Efficient Design) conference, which brought together presentations from all power-management efforts within IBM. Several external speakers were invited, including Dr. Graybill. The conference has been held annually since 2002. Multiple members of ARL attended and/or presented work at this conference.

## 9. Publications resulted from the project

### Book Chapters:

- [1] H. Aydin, R. Melhem, D. Mossé; "Periodic Reward-Based Scheduling and Its Application to Power-Aware Real-Time Systems", In *Handbook of Scheduling: Algorithms, Models and Performance Analysis*, CRC Press, 2003.
- [2] Power Aware Computing; Series in Computer Science, edited by R. Graybill and R. Melhem, *Kluwer Academic/Plenum Publishers*, May 2002
- [3] R. Melhem, N. AbouGhazaleh, H. Aydin and D. Mossé, "Power Management Points in Power-Aware Real-Time Systems", In *Power Aware Computing*, ed. by R. Graybill and R. Melhem, Plenum/Kluwer Publishers, 2002.
- [4] N. AbouGhazaleh, D. Mossé, B. Childers and R. Melhem; ``Toward The Placement of Power Management Points in Real Time Applications", In *Compilers and Operating Systems for Low Power*, Kluwer Academic Publishers, 2002.

### Journal publication:

- [5] P. Mejia-Alvarez, E. Levner and D. Mossé, "Adaptive Scheduling Server for Power-Aware Real-Time Tasks", *to appear in ACM Trans. on Embedded Computing Systems*.
- [6] C. Rusu, R. Melhem and D. Mossé, "Multi-version Scheduling in Rechargeable Energy-aware Real-time Systems", *to appear in Journal of Embedded Computing*.
- [7] D. Zhu, R. Melhem and D. Mossé, "Power Aware Scheduling for AND/OR Graphs in Real-Time Systems", *to appear in IEEE Trans. on Parallel and Distributed Systems*.
- [8] P. Mejia- Alvarez, R. Melhem, D. Mossé and H. Aydin, "An Incremental Server for Scheduling Overloaded Real-Time Systems", *IEEE Transactions on Computers*, Vol. 52, No. 10, 2003.
- [9] H. Aydin, R. Melhem, D. Mossé and P. Mejia-Alvarez, "Power-Aware Scheduling for Periodic Real-Time Tasks", *IEEE Trans. on Computers*, vol. 53, no. 5, pp. 584 - 600, 2004.
- [10] R. Melhem, D. Mossé and E.(Mootaz) Elnozahy, "The Interplay of Power Management and Fault Recovery in Real-Time Systems", *IEEE Trans. on Computers*, vol. 53, no. 2, pp. 217 - 231, 2004.
- [11] D. Zhu, R. Melhem, and B. Childers, "Scheduling with Dynamic Voltage/Speed Adjustment Using Slack Reclamation in Multi-Processor Real-Time Systems", *IEEE Trans. on Parallel & Distributed Systems*, vol. 14, no. 7, pp. 686 - 700, 2003.
- [12] C. Rusu, R. Melhem, D. Mossé, "Maximizing the System Value while Satisfying Time and Energy Constraints", *IBM Journal of R&D*, vol 47, no 5/6, 2003.
- [13] C. Rusu, R. Melhem, D. Mossé , "Maximizing Rewards for Real-Time Applications with Energy Constraints", *ACM TECS*, vol 2, no 4, 2003.

[14] P. Bohrer, M. Elnozahy, A. Gheith, C. Lefurgy, T. Nakra, J. Peterson, R. Rajamony, R. Rockhold, H. Shafi, R. Simpson, E. Speight, K. Sudeep, E. Van Hensbergen, and L. Zhang., "MAMBO - A Full System Simulator for the PowerPC Architecture". *ACM SIGMETRICS Performance Evaluation Review*, vol 3, no 4, pp. 8-12, 2004.

### **Conference and workshop publications:**

[15] D. Zhu, R. Melhem and D. Mossé; "Analysis of an Energy Efficient Optimistic TMR Scheme", International Conference on Parallel and Distributed Systems, *Newport Beach, CA, July 2004*.

[16] C. Rusu, R. Xu, R. Melhem, D. Mossé; "Energy-Efficient Policies for Request-Driven Soft Real-Time Systems", Euromicro Conference on Real-Time Systems, *Catania, July 2004*.

[17] N. AbouGhazaleh, P. Lanigan, S. Gobriel, D. Mossé and R. Melhem, "Dynamic Rate-Selection for Extending the Lifetime of Energy-Constrained Networks", to appear in Workshop on Energy Efficient Wireless Communication Networks, in conjunction with IPCCC 2004.

[18] S. Gobriel, R. Melhem and D. Mossé; "Unified Interference/Collision Analysis for Power Aware Adhoc Networks", IEEE INFOCOM 2004, *March 2004*.

[19] D. Zhu, D. Mossé and R. Melhem; "Multiple-Resource Periodic Scheduling Problem: how much fairness is necessary?", IEEE Real-Time Systems Symposium 2003 (RTSS'03), *Dec 2003*.

[20] N. AbouGhazaleh, B. Childers, D. Mossé, R. Melhem, and Matthew Craven "Energy Management for Real-Time Embedded Applications with Compiler Support", ACM SIGPLAN Languages, Compilers, and Tools for Embedded Systems (LCTES'03).

[21] N. AbouGhazaleh, D. Mossé, B. Childers, R. Melhem, and Matthew Craven "Collaborative Operating System and Compiler Power Management for Real-Time Applications", IEEE Real-Time Embedded Technology and Applications Symposium (RTAS), 2003.

[22] R. Mishra, N. Rastogi, D. Zhu, D. Mossé; R. Melhem; "Energy Aware Scheduling for Distributed Real-Time Systems", International Parallel and Distributed Processing Symposium (IPDPS'03), *Nice, France, Apr. 2003*.

[23] C. Rusu, R. Melhem, D. Mossé; "Multi-version Scheduling in Rechargeable Energy aware Real-time systems", Euromicro Conference on Real-Time Systems (ECRTS'03), *Porto, July 2003*.

[24] C. Rusu, R. Melhem, D. Mossé; "Maximizing the System Value while Satisfying Time and Energy Constraints", RTSS'02 (Real-Time Systems Symposium), *Austin, Dec 2002*.

[25] E. Elnozahy, R. Melhem, D. Mosse; "Energy-Efficient Duplex and TMR Real-Time Systems", RTSS'02 (Real-Time Systems Symposium), *Austin, Dec 2002*.

- [26] P. Mejia-Alvarez, E. Levner and D. Mossé; “Power-Optimized Scheduling Server for Real-Time Tasks”, Proceedings of the 8th Real-Time and Embedded Technology and Applications Symposium”, *San Jose, CA, September 2002*.
- [27] D. Zhu, N. AbouGhazaleh, D. Mossé and R. Melhem ; “Power Aware Scheduling for AND/OR Graphs in Multi-Processor Real-Time Systems”, ICPP’02 (International Conference on Parallel Processing), *Vancouver, B.C. Canada, Aug. 2002*.
- [28] P. Mejia-Alvarez, E. Levner and D. Mossé; “An Integrated Heuristic Approach to Power-Aware Real-Time Scheduling”, Workshop on Power Aware Computer Systems. *PACS 02. Feb, 2002. Springer Verlag’s Lecture Notes on Computer Science Series (LNCS 2325)*.
- [29] D. Zhu, R. Melhem, and B. Childers; “Scheduling with Dynamic Voltage/Speed Adjustment Using Slack Reclamation in Multi-Processor Real-Time Systems”, RTSS’01 (Real-Time Systems Symposium), *London, England, Dec 2001*.
- [30] H. Aydin, R. Melhem, D. Mossé, and Pedro Mejia Alvarez “Dynamic and Aggressive Scheduling Techniques for Power-Aware Real-Time Systems”, RTSS’01 (Real-Time Systems Symposium), *London, England, Dec 2001*.
- [31] N. AbouGhazaleh, D. Mossé, B. Childers and R. Melhem “Toward The Placement of Power Management Points in Real Time Applications”, COLP’01 (Workshop on Compilers and Operating Systems for Low Power), *Barcelona, Spain, 2001*.
- [32] H. Aydin, R. Melhem, D. Mossé and P. Mejia-Alvarez, “Determining Optimal Processor Speeds for Periodic Real Time Tasks with Different Power Characteristics”, ECRTS’01 (Euromicro Conference on Real-Time Systems), *Delft, Holland, 2001*.
- [33] B. Childers, H. Tang, and R. Melhem, “Adapting Processor Supply Voltage to Instruction-Level Parallelism”, Koolchips Workshop, in conjunction with MICRO-33, *Monterey, California, December 2000*.
- [34] D. Mossé, H. Aydin, B. Childers, and R. Melhem, “Compiler-Assisted Dynamic Power-Aware Scheduling for Real-Time Applications”, COLP (Workshop on Compiler and OS for Low Power), *Philadelphia, Oct 19, 2000*.
- [35] A. Allavena and D. Mossé, “Scheduling of Frame based Embedded Systems with Rechargeable Batteries”, Workshop on Power Management for Real-Time and Embedded Systems (in conjunction with RTAS 2001).
- [36] Pedro Mejia and D. Mossé; "Responsiveness Approach for Fault Recovery Operations in Real Time Systems.", Real Time Technology and Applications Symposium. *Vancouver, Canada, June 1999*.
- [37] M. Elnozahy, M. Kistler, and R. Rajamony, “Energy-Efficient Server Clusters”, In Proceedings of the Workshop on Power-Aware Computing Systems, February 2002.
- [38] M. Elnozahy, M. Kistler, and R. Rajamony, “Energy Conservation Policies for Web Servers”, In Proceedings of the 4th USENIX Symposium on Internet Technologies and Systems, Seattle, March, 2003.

[39] H. Shafi, P. Bohrer, J. Phelan, and C. Rusu, “Event-Based System Power Simulation”, In Proceedings of the IBM Austin Conference on Energy-Efficient Design (ACEED), March, 2002.